arXiv:astro-ph/0102340v1 20 Feb 2001

Submitted to ApJ

Computational Methods for Gravitational Lensing

Charles R. Keeton

Steward Observatory, University of Arizona,
933 N. Cherry Ave., Tucson, AZ 85721

ABSTRACT

Modern applications of strong gravitational lensing require the ability to use precise and varied observational data to constrain complex lens models. I discuss two sets of computational methods for lensing calculations. The first is a new algorithm for solving the lens equation for general mass distributions. This algorithm makes it possible to apply arbitrarily complicated models to observed lenses. The second is an evaluation of techniques for using observational data including positions, fluxes, and time delays of point-like images, as well as maps of extended images, to constrain models of strong lenses. The techniques presented here are implemented in a flexible and user-friendly software package called *gravlens*, which is made available to the community.

1. Introduction

Gravitational lensing is an important astrophysical tool because it directly probes mass (as opposed to luminosity) distributions, and because it brightens and enlarges the images of distant sources. The more than 60 known strong lenses[1] produced by galaxies constitute a *mass-selected* sample of galaxies at intermediate redshifts. Lensing provides precise mass measurements for these galaxies, and thereby offers a powerful probe of the physical properties of intermediate-redshift galaxies and the evolution of early-type galaxies in low-density environments (e.g., Keeton, Kochanek & Falco 1998; Kochanek et al. 2000). The lenses can be used to probe galaxy mass distributions (e.g., Kochanek 1991a; Rusin & Tegmark 2000; Rusin & Ma 2000; and references therein). They can also be used for direct measurements of the Hubble constant independent of the distance ladder (e.g., Koopmans & Fassnacht 1999; Witt, Mao & Keeton 2000; and references therein), to constrain the cosmological model (e.g., Falco, Kochanek & Muñoz 1998; Helbig et al. 1999; and references therein), and for detailed studies of the host galaxies of high-redshift quasars (e.g., Rix et al. 2000; Kochanek, Keeton & McLeod 2001a). Strong lenses produced by

[1]Strong lenses are systems with multiple images of a background source, and they are the focus of this paper. Weak lensing, or shape distortions without multiple imaging, also probes mass distributions but with different techniques. For recent reviews of weak lensing, see Mellier (1999) and Bartelmann & Schneider (2001).

clusters probe the radial mass distribution of clusters and reveal the clumpy galaxy distribution superimposed on the smooth cluster background, and thereby test models of structure formation in the cold dark matter paradigm (e.g., Tyson, Kochanski & Dell'Antonio 1998; Williams, Navarro & Bartelmann 1999; Shapiro & Iliev 2000).

In many of these eclectic lensing applications an essential step is fitting mass models to observed lenses. There are two key ingredients to modern lens models. The first is the ability to use the precise and varied observational data now available for many lenses. Optical and near-infrared astrometry with the Hubble Space Telescope achieves a precision of a few milli-arcseconds (e.g., Lehár et al. 2000). Radio astrometry from VLBI or VLBA maps can achieve a precision of 10 micro-arcseconds or better, and may resolve fine substructure in the images (e.g., Patnaik, Porcas & Browne 1995; Trotter, Winn & Hewitt 2000). Deep optical and especially near-infrared images often reveal extended structure due to the host galaxy of the source, and even some complete Einstein rings (e.g., Bernstein et al. 1997; Impey et al. 1998; Kochanek et al. 2001a). Photometry at many wavelengths and many epochs can reveal evidence for reddening and/or microlensing of the images (e.g., Gott et al. 1981; Falco et al. 1999; Wambsganss et al. 2000; Wozniak et al. 2000), and the monitoring can be used to measure time delays between the lensed images (e.g., Kundić et al. 1997a; Schechter et al. 1997). With the proper techniques, all of the different kinds of data can be used to constrain lens models.

The second important ingredient is the ability to study complex mass models. Reproducing qualitative features of a lens (the number and configuration of the images) can often be done with very simple models, but fitting the data quantitatively requires models that include detailed structure in the lens galaxy and its environment. For example, the models require an elliptical density distribution for early-type lens galaxies (e.g., Keeton & Kochanek 1997), or a thin disk and a rounder halo for spiral lens galaxies (e.g., Maller, Flores & Primack 1997), and perhaps even substructure in the galaxy (e.g., Mao & Schneider 1998; Bernstein & Fischer 1999). The models usually must include tidal perturbations from objects near the lens galaxy or along the line of sight (e.g., Young et al. 1981; Keeton, Kochanek & Seljak 1997; Witt & Mao 1997). While tidal perturbations are often approximated as an external shear for convenience, this simplification may be ruled out by data of sufficient quality (e.g., Impey et al. 1998). Several lenses are even found to have multiple galaxies that lie inside the Einstein ring and must be explicitly modeled (e.g., Koopmans & Fassnacht 1999; Rusin et al. 2000). The problem is that such complicated models can be hard to study due to the difficulty of solving the lens equation. For models with spherical or ellipsoidal symmetry all classes of solutions to the lens equation are known (see, e.g., Schneider, Ehlers & Falco 1992). For models without such symmetry, however, it may not even be clear what the maximum number of images is, much less how they are arranged. The need to use complex models to fit the data means that we need a general algorithm for solving the lens equation without requiring simplifying assumptions about symmetry.

This paper presents methods for modern lensing calculations, including a fully general algorithm for solving the lens equation that requires no assumptions about the symmetry of the

mass models, and techniques for handling a variety of observational data. The organization of the paper is as follows. Section 2 reviews the lens theory needed for most calculations. Section 3 presents the algorithm for solving the lens equation, while Section 4 discusses techniques for using observational data to constrain lens models. The methods presented here are implemented in a publicly-available software package called *gravlens*, which is described in Section 5. The software and documentation are available from the web site of the CfA/Arizona Space Telescope Lens Survey, at `http://cfa-www.harvard.edu/castles`.

2. Basic Lens Theory

Suppose a source at angular position $\mathbf{u}$ emits a light ray that passes a foreground mass distribution (the lens) with impact parameter $\mathbf{x}$ and gets deflected by the gravitational field of the lens. Compared to an undeflected ray, the deflected ray has a longer travel time because it has a longer geometric length and it passes through a gravitational potential well. Assuming that the lens is confined to a small fraction of the total path length, the extra light travel time is

$$\tau(\mathbf{x}) = \frac{1+z_l}{c}\,\frac{D_{ol}D_{os}}{D_{ls}}\left[\frac{1}{2}\,|\mathbf{x}-\mathbf{u}|^2 - \phi(\mathbf{x})\right], \tag{1}$$

where z_l is the redshift of the lens, and D_{ol}, D_{os}, and D_{ls} are angular diameter distances from the observer to the lens, from the observer to the source, and from the lens to the source, respectively. (See, e.g., Schneider et al. 1992 for a full discussion.) Also, ϕ is the two-dimensional gravitational potential of the lens,

$$\phi(\mathbf{x}) = \frac{1}{\pi}\int \frac{\Sigma(\mathbf{y})}{\Sigma_{cr}}\,\ln|\mathbf{x}-\mathbf{y}|\,d\mathbf{y}\,, \tag{2}$$

where Σ is the surface mass density of the lens, which is normalized by the critical surface density for lensing,

$$\Sigma_{cr} = \frac{c^2}{4\pi G}\,\frac{D_{os}}{D_{ol}D_{ls}}\,. \tag{3}$$

By Fermat's principle, images form at stationary points of the time delay surface, or at solutions of the equation

$$\mathbf{u} = \mathbf{x} - \nabla\phi(\mathbf{x})\,. \tag{4}$$

This is the general form of the gravitational lens equation (for a single lens plane).

In addition to producing a deflection, the lens also distorts and amplifies the image(s) in a manner described by the magnification tensor,

$$\mu \equiv \left(\frac{\partial\mathbf{u}}{\partial\mathbf{x}}\right)^{-1} = \begin{bmatrix} 1-\phi_{xx} & -\phi_{xy} \\ -\phi_{xy} & 1-\phi_{yy} \end{bmatrix}^{-1}, \tag{5}$$

where subscripts denote partial differentiation, $\phi_{ij} \equiv \partial^2\phi/\partial x_i \partial x_j$. Generically, a strong lens has one or more "critical curves" in the image plane along which $\det(\mu^{-1}) = 0$. The critical curves map

to "caustics" in the source plane, which are important because they are places where the number of images changes, and because a source near a caustic is highly amplified and distorted.

3. Solving the Lens Equation

Reading the lens equation from left to right — i.e., picking a source $\mathbf{u}$ and trying to find the images $\mathbf{x}_i$ — makes it hard to solve. Analytic solutions are often impossible because the deflection $\nabla\phi$ either involves transcendental functions or cannot be computed analytically. Numerical solutions are often difficult because there is no algorithm that is guaranteed to find all the roots of a two-dimensional equation (see Press et al. 1992). Hence, numerical techniques require independent knowledge of the number of images and rough guesses for their positions, but no local property of the lens equation gives this information. While the global caustic structure does characterize the number of images, it is hard to analyze except in relatively simple lens models with sufficient symmetry. The existence of compound lenses (e.g., Koopmans & Fassnacht 1999; Rusin et al. 2000) and lens galaxies with companions or satellites (e.g., Hogg & Blandford 1994; Kundić et al. 1997b; Tonry 1998) indicates the need for a general solver that does not require symmetry assumptions.

The key simplification is to reverse direction and read the lens equation from right to left, to view it as a mapping from the image plane to the source plane: it takes each image position $\mathbf{x}$ and maps it to a unique source position $\mathbf{u}(\mathbf{x}) = \mathbf{x} - \nabla\phi(\mathbf{x})$. This way of thinking leads to a straightforward solution to the problem that a numerical root finder needs to know the number and approximate locations of the images. Consider laying down some tiling of the image plane. The lens mapping takes each image plane tile I_j to a corresponding source plane tile S_j (by mapping vertices), and thus generates a tiling of the source plane that covers every point with at least one tile. Figure 1 depicts a sample tiling. The tiling actually contains all of the information we need to solve the lens equation. There is some set of tiles $(S_j, S_k, \ldots)$ that cover any particular source position, and the number of covering tiles gives the number of images of that source. Furthermore, the corresponding image plane tiles $(I_j, I_k, \ldots)$ bound the image positions.[2] In other words, the tiling easily reveals both the number of images and their approximate positions, which can be efficiently refined with a numerical root finder.

The lens mapping incorporates all of the properties of the lens equation, such as the folding and stretching that determine the number of images and their distortions (see Figure 1). The tiling automatically includes all of the features of the lens mapping (albeit at finite resolution), without requiring any assumptions about symmetry in the lensing mass. As a result, the tiling algorithm provides a fully general method for solving the lens equation to analyze the lensing properties of arbitrary mass distributions. Blandford & Kochanek (1987) and Kochanek & Blandford (1987) have used a similar technique for calculations of lens statistics. More generally, the tiling algorithm

[2]If a source is covered by source tile S_j, then an image must lie within the corresponding image plane tile I_j.

can be used to find the critical curves and caustics and to solve the lens equation for arbitrarily complicated mass distributions, such as the random collection of five galaxies shown in Figure 2.

Implementing the tiling algorithm reveals three important details. First, the tiles should be triangles, because triangles are the only polygons that are guaranteed to remain convex regardless of how they are deformed by the mapping. Also, triangular tiles make it easy to determine whether a given source point $\mathbf{u}$ is covered by a tile. Let $\mathbf{u}_i$ $(i = 1, 2, 3)$ be the vertices of the tile, and let $\delta\mathbf{u}_i = \mathbf{u}_i - \mathbf{u}$ be the vector from the source point to vertex i; the source is contained within the triangle if the three cross products $\delta\mathbf{u}_1 \times \delta\mathbf{u}_2$, $\delta\mathbf{u}_2 \times \delta\mathbf{u}_3$, and $\delta\mathbf{u}_3 \times \delta\mathbf{u}_1$ all have the same sign. In practice, it may be convenient to use quadrilateral tiles (see Figure 1) and divide each one into two triangles.

Second, we must consider the resolution of the tiling. Good resolution is important near the critical curves in order to resolve the folding (see Figures 1 and 3), but to avoid unnecessary calculations we want an adaptive algorithm that uses high resolution only where necessary. Fortunately there is a simple local criterion that identifies such regions. If the scalar magnification, $\det(\mu)$, changes sign across an image plane tile then that tile contains a critical curve. To increase the resolution, the tile can be broken into an array of sub-tiles. This sub-tiling can then be recursively repeated to obtain even better resolution. Figure 3 illustrates a tiling that has three levels of sub-tiling near the critical curves. In lens models with more than one galaxy, recursive sub-tiling can also be used near the additional galaxies to resolve their critical curves. The sub-tiling offers the additional benefit of putting tight bounds on the locations of the critical curves, which can then be refined with a numerical root finder.

Finally, the sub-tiling introduces one additional technicality. Consider placing the vertex of a daughter tile on the edge of its parent tile. Because a straight line in the image plane may map to a curve in the source plane, in the source plane the daughter vertex map *not* lie on the edge of the parent tile. As a result, the daughter tiles may not cover exactly the same area as the parent tile; the result can be gaps or overlaps in the tiling, as illustrated in Figure 4. This problem occurs only at places where the resolution changes. (If adjacent tiles are both sub-tiled, the gap in one sub-tiling is exactly offset by the overlap in the other; see point C in Figure 4.) The problem can be easily corrected by using 2×2 sub-tiling so that each gap or overlap is a triangle that can be explicitly examined.

4. Strategies for Modeling Strong Lenses

Given a general lens solver, an important application is fitting models to observed lenses. Lens observations now offer a wide range of high-precision data, and this section presents a variety of techniques for using the data to constrain models. For lenses with point-like images, the positions, fluxes, and time delays of the images can be used with least-squares fitting methods, and Sections 4.1–4.3 give definitions of the appropriate goodness of fit statistics. It may be possible to use the

image positions to solve for certain model parameters analytically, as shown in Section 4.4. Finally, for lenses with extended images like jets, arcs, and Einstein rings, Section 4.5 briefly reviews techniques for using them as constraints.

4.1. Image positions

With high-resolution optical or radio observations, it is usually reasonable to assume that the astrometric uncertainties for different point images are independent and Gaussian. The χ^2 term for the image positions, evaluated in the *image plane*, is then

$$\chi^2_{img} = \sum_i \delta\mathbf{x}_i^T \cdot S_i^{-1} \cdot \delta\mathbf{x}_i \,, \tag{6}$$

$$\delta\mathbf{x}_i = \mathbf{x}_{obs,i} - \mathbf{x}_{mod,i} \,, \tag{7}$$

where the sum extends over all images, and $\mathbf{x}_{obs,i}$ and $\mathbf{x}_{mod,i}$ are the observed and modeled positions of image i. The astrometric uncertainties for image i are described by the covariance matrix

$$S_i = R_i^T \begin{bmatrix} \sigma_{1,i}^2 & 0 \\ 0 & \sigma_{2,i}^2 \end{bmatrix} R_i \,, \tag{8}$$

$$R_i = \begin{bmatrix} -\sin\theta_{\sigma,i} & \cos\theta_{\sigma,i} \\ -\cos\theta_{\sigma,i} & -\sin\theta_{\sigma,i} \end{bmatrix} , \tag{9}$$

where the error ellipse has semi-major axis $\sigma_{1,i}$, semi-minor axis $\sigma_{2,i}$, and position angle $\theta_{\sigma,i}$ (measured East of North).

There is an alternate position χ^2 that is evaluated in the *source plane* (e.g., Kayser et al. 1990; Kochanek 1991a),

$$\chi^2_{src} = \sum_i \delta\mathbf{u}_i^T \cdot \mu_i^T \cdot S_i^{-1} \cdot \mu_i \cdot \delta\mathbf{u}_i \,, \tag{10}$$

$$\delta\mathbf{u}_i = \mathbf{u}_{obs,i} - \mathbf{u}_{mod} \,, \tag{11}$$

where $\mathbf{u}_{obs,i} = \mathbf{x}_{obs_i} - \nabla\phi(\mathbf{x}_{obs_i})$ is the source corresponding observed image i, $\mathbf{u}_{mod}$ is the model source position, μ_i is the magnification tensor for image i, and the covariance matrix S_i is as above. The factors of the magnification tensor are inserted because if the source plane deviation $\delta\mathbf{u}_i$ is small enough that the magnification is nearly constant, then $\mu_i \cdot \delta\mathbf{u}_i \approx \delta\mathbf{x}_i$ yields an approximate image plane deviation. In other words, χ^2_{src} is an approximate version of χ^2_{img}. The approximation can be useful because it not only avoids the need to solve the lens equation but also allows the best model source position to be found analytically:

$$\mathbf{u}_{mod} = A^{-1} \cdot \mathbf{b} \,, \tag{12}$$

$$\text{where} \quad A = \sum_i \mu_i^T \cdot S_i^{-1} \cdot \mu_i \,, \tag{13}$$

$$\mathbf{b} = \sum_i \mu_i^T \cdot S_i^{-1} \cdot \mu_i \cdot \mathbf{u}_{obs,i} \,, \tag{14}$$

which is straightforward to evaluate because the matrices are 2×2. However, the approximation inherent in χ^2_{src} is somewhat undesirable for two reasons. First, strictly speaking χ^2_{src} is an accurate representation of χ^2_{img} only if $\delta\mathbf{u}$ is small; in other words, χ^2_{src} should properly be used only if a good model is already known, not in an initial search for a good model. Second, χ^2_{src} is not computed with directly observable quantities. Despite these limitations, χ^2_{src} is qualitatively correct in the sense that for poor models it gives a large value of χ^2 and thus indicates a poor fit. In summary, the fact that χ^2_{src} is computationally fast and at least qualitatively accurate means that it can be useful for initial modeling to find the appropriate region of parameter space to explore. The problems with the approximation, however, mean that it is preferable to use χ^2_{img} for refining models to find the best-fit model and the range of models consistent with the data.

4.2. Image fluxes

The (relative) fluxes of point images can provide useful model constraints, but they must be used carefully. The problem is that the statistical errorbars may underestimate the actual uncertainties. The fluxes, especially at optical wavelengths, may be affected by microlensing, small-scale structure in the lens galaxy, and differential reddening by dust in the lens galaxy (e.g., Mao & Schneider 1998; Falco et al. 1999; Wozniak et al. 2000). They can also be affected by source variability and time delays. If these systematic effects are not understood and corrected, the strength of the flux constraints must not be overemphasized. A simple and common way of using the fluxes conservatively is to retain the assumption of Gaussian errors for simplicity, but to inflate the errorbars to represent an estimate of the systematic uncertainties (e.g., Koopmans & Fassnacht 1999; Cohn et al. 2000).

If the photometric errors for the various images are independent and Gaussian, the χ^2 for the fluxes is

$$\chi^2_{flux} = \sum_i \frac{(f_i - M_i f_{src})^2}{\sigma^2_{f,i}} \ , \tag{15}$$

where the observed flux of image i is $f_i \pm \sigma_{f,i}$, and the model gives the magnification $M_i = |\det(\mu_i)|$ of image i and the intrinsic flux f_{src} of the source. The best-fit source flux can be found analytically,

$$f_{src} = \frac{\sum_i f_i \, M_i/\sigma^2_{f,i}}{\sum_i M_i^2/\sigma^2_{f,i}} \ . \tag{16}$$

If it is preferable to express the photometry in magnitudes (rather than physical fluxes), the χ^2 term is

$$\chi^2_{mag} = \sum_i \frac{(m_i + 2.5 \log M_i - m_{src})^2}{\sigma^2_{m,i}} \ , \tag{17}$$

where the observed magnitude of image i is $m_i \pm \sigma_{m,i}$. The best-fit source magnitude is then

$$m_{src} = \frac{\sum_i (m_i + 2.5 \log M_i)/\sigma^2_{m,i}}{\sum_i 1/\sigma^2_{m,i}} \ . \tag{18}$$

In all of these equations the units are arbitrary; they can be absolute fluxes or magnitudes, or they can be relative fluxes or magnitudes referenced to one of the images.

The assumption that the uncertainties in the (relative) photometry are independent holds only if the images do not have overlapping point spread functions. For lenses that are poorly resolved, deconvolution may be able to separate the various images (e.g., Courbin et al. 1997), but may (in principle) yield correlated photometric errors. The flux χ^2 is easily generalized to handle correlated errors,

$$\chi^2_{flux} = \sum_{ij} (S_f^{-1})_{ij} (f_i - M_i f_{src})(f_j - M_j f_{src}) \,, \tag{19}$$

where S_f is the covariance matrix of the fluxes, and both the i and j sums extend over all images. The best-fit source flux is then

$$f_{src} = \frac{\sum_{ij} (S_f^{-1})_{ij} M_i f_j}{\sum_{ij} (S_f)^{-1}_{ij} M_i M_j} \,. \tag{20}$$

If it is preferable to use magnitudes, the χ^2 term is

$$\chi^2_{mag} = \sum_{ij} (S_m^{-1})_{ij} (m_i + 2.5 \log M_i - m_{src})(m_j + 2.5 \log M_j - m_{src}) \,, \tag{21}$$

where S_m is the covariance matrix of the magnitudes. In this case the best-fit source magnitude is

$$m_{src} = \frac{\sum_{ij} (S_m^{-1})_{ij} (m_j + 2.5 \log M_j)}{\sum_{ij} (S_m^{-1})_{ij}} \,. \tag{22}$$

4.3. Time delays

By measuring light curves for the images and comparing them to each other, it may be possible to measure the time delay(s) induced by lensing. In a lens with N images there are up to $N-1$ independent time delays. They can be used with lens models to determine the Hubble constant H_0 (e.g., Refsdal 1964), and to provide up to $N-2$ additional constraints on the models. To understand how time delays are used, let $\tau_{obs,i}$ be the time by which the light curve of image i lags the light curve of the leading image. The prediction from a lens model for the delay can be decomposed as follows (see eq. 1):

$$\tau_{mod,i} = h^{-1} \, t_0 \, \bar{\tau}_{mod,i} \,, \tag{23}$$

$$t_0 = \frac{1+z_l}{c} \frac{D_{ol} D_{os}}{D_{ls}} \,, \qquad \text{(computed for } h = 1) \tag{24}$$

$$\bar{\tau}_{mod,i} = \frac{1}{2} \Big[|\mathbf{x}_i - \mathbf{u}|^2 - |\mathbf{x}_{lead} - \mathbf{u}|^2 \Big] - \Big[\phi(\mathbf{x}_i) - \phi(\mathbf{x}_{lead}) \Big] , \tag{25}$$

where $h = H_0/(100 \text{ km s}^{-1} \text{ Mpc}^{-1})$ and the subscript "lead" indicates the leading image. Eq. (23) shows that the model time delay includes a factor that depends only on cosmology (t_0) and a factor

that depends only on the lens model ($\bar{\tau}_{mod,i}$); it also shows the explicit dependence on the Hubble parameter h.

If the errors in the time delays are independent and Gaussian, the χ^2 contribution is

$$\chi^2_{tdel} = \sum_i \frac{\left(\tau_{obs,i} - h^{-1}\, t_0\, \bar{\tau}_{mod,i}\right)^2}{\sigma^2_{t,i}} \ . \tag{26}$$

To find the best-fit value of h together with confidence intervals, one approach is to trace out χ^2 versus h. An alternative approach is to let h be free to adopt the best-fit value, although perhaps with a prior range $h_{prior} \pm \sigma_{h,prior}$ specified to ensure that the value is reasonable. (To let h be unrestricted, i.e. with no prior range specified, simply use a large value for $\sigma_{h,prior}$.) The χ^2 contribution for the prior range is

$$\chi^2_h = \frac{\left(h^{-1} - h^{-1}_{prior}\right)^2}{\sigma^2_h} \ , \qquad \text{where} \qquad \sigma_h = \frac{\sigma_{h,prior}}{h^2_{prior}} \ . \tag{27}$$

Note that this χ^2 contribution is written in terms of h^{-1} rather than h, and σ_h is the uncertainty in h^{-1}, while $\sigma_{h,prior}$ is the uncertainty in h itself. This is done for convenience, making it possible to solve analytically for the best-fit value of h (the value that minimizes $\chi^2_{tdel} + \chi^2_h$),

$$h = \left[\frac{1}{\sigma^2_h} + \sum_i \frac{t_0^2\, \bar{\tau}^2_{mod,i}}{\sigma^2_{t,i}}\right] \Bigg/ \left[\frac{1}{h_{prior}\, \sigma^2_h} + \sum_i \frac{t_0\, \bar{\tau}_{mod,i}\, \tau_{obs,i}}{\sigma^2_{t,i}}\right] . \tag{28}$$

With a single measured time delay τ_{obs} and no prior assumption about H_0, eq. (28) is equivalent to $h = t_0\, \bar{\tau}_{mod}/\tau_{obs}$. If there are M time delays, they determine a value for h and also provide $M - 1$ constraints on lens models.

A potential problem with this analysis is that the time delay errors may be very non-Gaussian. The light curves used to measure time delays usually have irregular sampling and may have substantial non-Gaussian "noise" from phenomena such as microlensing (e.g., Thomson & Schild 1997; Fassnacht et al. 1999). A relatively simple test can be done to examine how conclusions about the Hubble constant are affected by non-Gaussian errors. One way to estimate the error distribution for the time delays is to perform Monte Carlo simulations with the light curve data; the simulations yield a set of time delays whose distribution represents the uncertainties. Using all of the simulated time delays to determine the Hubble constant produces a distribution of H_0 values that incorporates all of the time delay errors, including not only non-Gaussian properties but also any correlations that might exist. Fassnacht et al. (1999) and Koopmans & Fassnacht (1999) have performed this test using the three time delays measured for the 4-image lens B 1608+656. Assuming independent non-Gaussian errors, i.e. using eq. (26), they find $H_0 = 59^{+7}_{-6}$ km s^{-1} Mpc^{-1}; by contrast, when they use the full error distribution from Monte Carlo simulations they find $H_0 = 59^{+8}_{-7}$ km s^{-1} Mpc^{-1}. (Both sets of errorbars encompass the 95% confidence range.) While this single test cannot rule out the importance of non-Gaussian time delay errors in all lensing determinations of H_0, it does suggest that their effects may not be dramatic.

4.4. Linear parameters and constraints

In parametric lens models, the lensing potential, deflection, and magnification are often complicated functions of the parameters. Sometimes, however, there are parameters that enter as simple linear coefficients. For example, the potential of a softened isothermal sphere lens is

$$\phi(r) = b\left[\sqrt{s^2+r^2} - s - s\,\ln\left(\frac{s+\sqrt{s^2+r^2}}{2s}\right)\right]. \tag{29}$$

The core radius s is a non-linear parameter because it appears in the arguments of transcendental functions. The mass parameter b, by contrast, is a linear parameter because it appears as a simple multiplicative coefficient. For non-linear parameters, the only way to find the best-fit values is to use a numerical algorithm to find minima of the multi-dimensional χ^2 function. For linear parameters, however, it may be possible to compute their best-fit values analytically, thus reducing the number of parameters that must be varied in the numerical optimization.

In principle, position, flux, and time delay constraints can all be used to solve for linear parameters. In practice, however, the flux and time delay data are usually too unreliable to be used this way. Besides, the image positions often provide more than enough constraints to solve for all of the linear parameters in a model, so this section discusses only position data. Linear techniques have been used by, e.g., Kochanek (1991b), Bernstein & Fischer (1999), and Keeton et al. (2000a), but this section offers a more general discussion.

Consider N images $\{\mathbf{x}_i\}_{i=1}^N$ of a given source. Requiring that the images be fit *exactly* gives N vector equations,

$$\mathbf{u} = \mathbf{x}_i - \nabla\phi\left(\mathbf{x}_i\right), \qquad (i = 1, \ldots, N). \tag{30}$$

Eliminating the unobservable source position, this is equivalent to $N-1$ vector equations

$$\mathbf{x}_1 - \nabla\phi\left(\mathbf{x}_1\right) = \mathbf{x}_j - \nabla\phi\left(\mathbf{x}_j\right), \qquad (j = 2, \ldots, N). \tag{31}$$

Because the vector space is two-dimensional, this amounts to $2(N-1)$ constraint equations.

Now suppose that some of the parameters in the lens model are linear, i.e. they enter as simple multiplicative coefficients in the lensing potential (like b in eq. 29). Let the M linear parameters be $\mathbf{p} = \{p_\alpha\}_{\alpha=1}^M$ and the M' non-linear parameters be $\mathbf{q} = \{q_\beta\}_{\beta=1}^{M'}$. The definition of linear parameters is that the lensing potential can be decomposed as

$$\phi(\mathbf{x};\mathbf{p},\mathbf{q}) = \sum_{\alpha=1}^{M} p_\alpha\,\varphi_\alpha(\mathbf{x};\mathbf{q}) \tag{32}$$

for some set of functions $\{\varphi_\alpha(\mathbf{x};\mathbf{q})\}_{\alpha=1}^M$ that are functions of position and are parametrized by only the non-linear parameters $\mathbf{q}$.

If the potential has the form eq. (32), then the constraints in eq. (31) become

$$\mathbf{x}_1 - \sum_{\alpha=1}^{M} p_\alpha\,\nabla\varphi_\alpha(\mathbf{x}_1;\mathbf{q}) = \mathbf{x}_j - \sum_{\alpha=1}^{M} p_\alpha\,\nabla\varphi_\alpha(\mathbf{x}_j;\mathbf{q}), \qquad (j = 2, \ldots, N). \tag{33}$$

If the number of constraints equals the number of linear parameters, i.e. $M = 2(N-1)$, then eq. (33) is a square matrix equation that can be solved to find the set of linear parameters $\mathbf{p}$ that yield an exact fit to the image positions. Because the constraint equations depend on the non-linear parameters $\mathbf{q}$, solving eq. (33) actually gives the linear parameters as functions of the non-linear parameters, $\mathbf{p}(\mathbf{q})$. Formally, this amounts to solving the matrix equation

$$A(\mathbf{q}) \cdot \mathbf{p}(\mathbf{q}) = \mathbf{b}\,, \tag{34}$$

where the matrix A has the components $A_{i\alpha}$ and the vector $\mathbf{b}$ has the components b_i such that

$$A_{(2j-3)(\alpha)} = \frac{\partial}{\partial x}\varphi_\alpha(\mathbf{x}_j;\mathbf{q}) - \frac{\partial}{\partial x}\varphi_\alpha(\mathbf{x}_1;\mathbf{q})\,, \qquad (j=2,\ldots,N) \tag{35}$$

$$A_{(2j-2)(\alpha)} = \frac{\partial}{\partial y}\varphi_\alpha(\mathbf{x}_j;\mathbf{q}) - \frac{\partial}{\partial y}\varphi_\alpha(\mathbf{x}_1;\mathbf{q})\,, \tag{36}$$

$$b_{(2j-3)} = x_j - x_1\,, \tag{37}$$

$$b_{(2j-2)} = y_j - y_1\,. \tag{38}$$

We can generalize this technique to sets of images associated with different sources. Suppose there are N images $\{\mathbf{x}_i\}_{i=1}^{N}$ of one source, and N' images $\{\mathbf{x}'_i\}_{i=1}^{N'}$ of a second source. Then the full set of constraints is

$$\mathbf{x}_1 - \nabla\phi\left(\mathbf{x}_1\right) = \mathbf{x}_j - \nabla\phi\left(\mathbf{x}_j\right)\,, \qquad (j=2,\ldots,N)\,, \tag{39}$$

$$\mathbf{x}'_1 - \nabla\phi\left(\mathbf{x}'_1\right) = \mathbf{x}'_k - \nabla\phi\left(\mathbf{x}'_k\right)\,, \qquad (k=2,\ldots,N')\,, \tag{40}$$

(compare eq. 31). Again finding the proper number of linear parameters, now $M = 2(N+N'-2)$, leads to a solvable matrix equation similar to eq. (34), where the generalization of the matrix A and vector $\mathbf{b}$ is straightforward: simply insert the entries for the second set of images $\{\mathbf{x}'_i\}_{i=1}^{N'}$ as additional rows in A and $\mathbf{b}$. The generalization to further sets of images is similar.

When used properly, linear techniques can speed up model searches by eliminating some parameters and by automatically identifying models that match at least some of the data. These features are especially helpful when working with high-precision data, because good models are likely to be confined to a narrow range of parameter space, so the χ^2 surface is likely to be a high plateau cut by narrow valleys in a way that frustrates multi-dimensional optimization algorithms (see Press et al. 1992). Without linear techniques, the optimization routine may wander around and miss the valleys; but with linear techniques, the optimization routine is directly guided to the valleys that need to be explored. Two important points must guide the use of linear techniques, however. The first is that number of linear parameters and constraints must match. For example, the relative positions in a 2-image lens provide two constraints, so they require two linear model parameters to be used with linear techniques. The relative positions in a 4-image lens provide six constraints, so they require six linear parameters.

The second important point is that using linear techniques forces the model to fit the data *exactly*, which is much stronger than asking the model to reproduce the data within the errorbars.

Consider a lens (call it A) produced by a smooth mass distribution. Also consider a second lens (call it A') where the same smooth mass distribution is perturbed by small-scale structure such as an isophote twist or a lumpy mass distribution. The small-scale structure can cause the image positions in lens A' to differ from those in lens A by up to $\sim$1 milli-arcsecond, even for the same source position (see Mao & Schneider 1998). Suppose the observational uncertainties are $\gtrsim$1 mas, so a simple smooth lens model yields a perfect fit to lens A and a worse but statistically acceptable fit to lens A'. Now if linear techniques are applied to lens A', the smooth model will be forced to adjust large-scale mass components to try to fit the data exactly. The model may be pushed into a strange corner of parameter space, or it may simply report that it cannot fit the data. Either way, the modeler will be led to believe that there are significant differences between the lens galaxies in A and A', when in fact the only difference is in small-scale mass components. A similar argument holds even when linear techniques are not used, if the observational uncertainties are much smaller than 1 mas.

The mistake here is the use of an oversimplified lens model with high-precision data. The lesson is that because of the possibility of perturbations from substructure, simple smooth lens models should not be asked to fit data to better than $\sim$1 mas precision. In other words, when using high-precision data in general, and specifically when using linear techniques, it is very important to allow complexity in the lens model. As an example, Keeton et al. (2000a; also see Kochanek 1991b; Bernstein & Fischer 1999) show that allowing physically-motivated substructure in the models is essential for robustly measuring the large-scale properties of the lens galaxy in the lens Q 0957+561.

4.5. Extended images

Extended images like arcs or Einstein rings can dramatically increase the number of constraints on models (compared with the constraints from point images). Extended images must be treated with some care, though, and several techniques have been developed to deal with them. Some radio images that appear point-like in low- or moderate-resolution radio maps can be resolved into sub-components by high-resolution VLBI or VLBA maps. The sub-components can be used as multiple sets of point-like images that have high-precision astrometry. See Trotter et al. (2000) for an excellent recent example of VLBI data and modeling. For smooth extended images like arcs and Einstein rings, several "exhaustive" algorithms have been developed to analyze the full observed map while allowing for the fact that the intrinsic source is unknown (e.g., Kochanek et al. 1989; Kochanek & Narayan 1992; Wallington, Kochanek & Narayan 1996; Tyson et al. 1998; Keeton et al. 2000a). These techniques extract the maximum amount of information from the images, but they are computationally expensive. Kochanek et al. (2001ab) have developed two alternate techniques — "ring fitting" for optical or infrared Einstein rings, and "curve fitting" for arcs — that extract less information from the images but are computationally fast. At least one of these techniques for extended images should be appropriate for any particular lens, and together with the various techniques for handling point-like images they make it possible to use virtually any kind of

observational data to constrain lens models.

5. A Software Package

The techniques presented in this paper have been used to create a versatile software package called *gravlens* that can be used for a wide variety of lensing calculations. The code has a command-driven interface that is easy to use, and it comes with a detailed manual. The software and documentation can be downloaded from the web site for the CfA/Arizona Space Telescope Lens Survey, at `http://cfa-www.harvard.edu/castles`.

The software is actually a package of two applications. One application, called *gravlens*, provides capabilities for basic lensing calculations. It combines the tiling algorithm from §3 with an extensive catalog of circular and elliptical mass models that can be arranged into arbitrarily complex composite models. Thus, it can be used to study the general lensing properties of virtually any mass distribution. It has been used by Keeton & Kochanek (1998), Keeton, Mao & Witt (2000b), and Rusin et al. (2000) to analyze the caustic structures in complex lens models.

The second application, called *lensmodel*, begins with all of the capabilities of *gravlens* and adds many features for modeling observed lenses (as summarized in Table 1). The *lensmodel* software allows any or all of the observational constraints discussed in §4, so it can be used with complex lenses that include point images of multiple sources and/or extended images like arcs or rings. It allows arbitrarily complicated mass models to be fit to the data, so it can be used for essentially any lens in any environment. It gives the user complete control over the behavior of the model parameters; for example, certain parameters can be varied while others are held fixed, constraints can be placed on parameter ranges, and explicit relations can even be imposed between the parameters. The parameter controls allow a thorough understanding of any covariances, degeneracies, or other systematic uncertainties in the models and the conclusions drawn from them. In other words, *lensmodel* offers the ability to fit complex mass models to sophisticated data, and it provides the control needed to fully analyze the range of models consistent with the data. The *lensmodel* software has been used to study models of many observed lenses (Keeton & Kochanek 1997; Keeton et al. 1998, 2000a; Cohn et al. 2000; Kochanek et al. 2000b; Lehár et al. 2000; Ros et al. 2000; and Rusin et al. 2000).

Acknowledgements: The techniques and software presented here have been under development for several years, and during that time many people have contributed to the work. Chris Kochanek has offered a tremendous number of ideas, and has thoroughly tested each feature as it was added to the code. Joanne Cohn, Jose Muñoz, David Rusin, Brian McLeod, and Joseph Lehár have repeatedly put the code through its paces, uncovering bugs of many shapes and sizes and offering helpful suggestions for the code and documentation. Finally, an anonymous referee gave an insightful critique that improved the presentation of this paper. Support for this work has been

provided by ONR-NDSEG grant N00014-93-I-0774, NSF grant AST-9407122, NASA ATP grant NAG5-4062, and Steward Observatory.

REFERENCES

Bartelmann, M., & Schneider, P. 2001, Physics Reports, 340, 291

Bernstein, G., Fischer, P., Tyson, J. A., & Rhee, G. 1997, ApJ, 483, L79

Bernstein, G., & Fischer, P. 1999, AJ, 118, 14

Blandford, R., & Kochanek, C. S. 1987, ApJ, 321, 658

Cohn, J. D., Kochanek, C. S., McLeod, B. A., & Keeton, C. R. 2000, preprint (astro-ph/0008390)

Courbin, F., Magain, P., Keeton, C. R., Kochanek, C. S., Vanderriest, C., Jaunsen, A. O., & Hjorth, J. 1997, A&A, 324, L1

Falco, E. E., Kochanek, C. S., & Muñoz, J. A. 1998, ApJ, 494 47

Falco, E. E., Impey, C. D., Kochanek, C. S., Lehár, J., McLeod, B. A., Rix, H.-W., Keeton, C. R., Muñoz, J. A., & Peng, C. Y. 1999, ApJ, 523, 617

Gott, J. R. 1981, ApJ, 243, 140

Helbig, P., Marlow, D., Quast, R., Wilkinson, P. N., Browne, I. W. A., & Koopmans, L. V. E. 1999, A&AS, 136, 297

Hogg, D. W., & Blandford, R. D. 1994, MNRAS, 268, 889

Impey, C. D., Falco, E. E., Kochanek, C. S., Lehár, J., McLeod, B. A., Rix, H.-W., Peng, C. Y., and Keeton, C. R. 1998, ApJ, 509, 551

Kayser, R., Surdej, J., Condon, J. J., Kellerman, K. I., Magain, P., Remy, M., & Smette, A. 1990, ApJ, 364, 15

Keeton, C. R., Kochanek, C. S., & Seljak, U. 1997, ApJ, 482, 604

Keeton, C. R., & Kochanek, C. S. 1997, ApJ, 487, 42

Keeton, C. R., & Kochanek, C. S. 1998, ApJ, 495, 157

Keeton, C. R., Kochanek, C. S., & Falco, E. E. 1998, ApJ, 509, 561

Keeton, C. R., Falco, E. E., Impey, C. D., Kochanek, C. S., Lehár, J., McLeod, B. A., Rix, H.-W., Muñoz, J. A., & Peng, C. Y. 2000a, ApJ, 542, 74

Keeton, C. R., Mao, S., & Witt, H. J. 2000b, ApJ, 537, 697

Kochanek, C. S., & Blandford, R. 1987, ApJ, 321, 658

Kochanek, C. S., Blandford, R. D., Lawrence, C. R., & Narayan, R. 1989, MNRAS, 238, 43

Kochanek, C. S. 1991a, ApJ, 373, 354

Kochanek, C. S. 1991b, ApJ, 382, 58

Kochanek, C. S., & Narayan, R. 1992, ApJ, 401, 461

Kochanek, C. S., Falco, E. E., Impey, C. D., Lehár, J., McLeod, B. A., Rix, H.-W., Keeton, C. R, Muñoz, J. A., & Peng, C. Y. 2000a, ApJ, 543, 131

Kochanek, C. S., Keeton, C. R., & McLeod, B. A. 2001a, ApJ, 547, 50

Kochanek, C. S., et al. 2001b, in preparation

Koopmans, L. V. E., & Fassnacht, C. D. 1999, ApJ, 527, 513

Kundić, T., et al. 1997a, ApJ, 482, 75

Kundić, T., Hogg, D. W., Blandford, R. D., Cohen, J. G., & Lubin, L. M. 1997b, AJ, 114, 2276

Lehár, J., Falco, E. E., Kochanek, C. S., McLeod, B. A., Muñoz, J. A., Impey, C. D., Rix, H.-W., Keeton, C. R., & Peng, C. Y. 2000, ApJ, 536, 584

Maller, A. H., Flores, R. A., & Primack, J. R. 1997, ApJ, 486, 681

Mao, S., & Schneider, P. 1998, MNRAS, 295, 587

Mellier, Y. 1999, ARA&A, 37, 127

Patnaik, A. R., Porcas, R. W., & Browne, I. W. A. 1995, MNRAS, 274, L5

Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, Numerical Recipes in C: The Art of Scientific Computing, Second Edition (New York: Cambridge Univ. Press)

Refsdal, S. 1964, MNRAS, 128, 307

Rix, H.-W., et al. 2000, ApJ, submitted

Ros, E., Guirado, J. C., Marcaide, J. M., Pérez-Torres, M. A., Falco, E. E., Muñoz, J. A., Alberdi, A., & Lara, L. 2000, A&A, 362, 845

Rusin, D., et al. 2000, preprint (astro-ph/0011505)

Rusin, D., & Ma, C.-P. 2000, preprint (astro-ph/0009079)

Rusin, D., & Tegmark, M. 2000, preprint (astro-ph/0008329)

Schechter, P. L., et al. 1997, ApJ, 475, L85

Schneider, P., Ehlers, J., & Falco, E. E. 1992, Gravitational Lenses (New York: Springer)

Shapiro, P. R., & Iliev, I. T. 2000, ApJ, 542, L1

Thomson, D. J., & Schild, R. 1997, in Applications of Time Series Analysis in Astronomy and Meteorology (New York: Chapman and Hall), ed. T. Subba Rao, M. B. Priestley & O. Lessi, p. 187

Tonry, J. L. 1998, AJ, 115, 1

Trotter, C. S., Winn, J. N., & Hewitt, J. N. 2000, ApJ, 535, 671

Tyson, J. A., Kochanski, G. P., & Dell'Antonio, I. P. 1998, ApJ, 498, L107

Wallington, S., Kochanek, C. S., & Narayan, R. 1996, ApJ, 465, 64

Wambsganss, J., Schmidt, R. W., Colley, W., Kundic, T., & Turner, E. L. 2000, A&A, 362, L37

Williams, L. L. R., Navarro, J. F., & Bartelmann, M. 1999, ApJ, 527, 535

Witt, H. J., & Mao, S. 1997, MNRAS, 291, 211

Witt, H. J., Mao, S., & Keeton, C. R. 2000, ApJ, 544, 198

Wozniak, P. R., Udalski, A., Szymanski, M., Kubiak, M., Pietrzynski, G., Soszynski, I., & Zebrun, K. 2000, ApJ, 540, L65

Young, P., Gunn, J. E., Kristian, J., Oke, J. B., & Westphal, J. A. 1981, ApJ, 244, 736

This preprint was prepared with the AAS LaTeX macros v5.0.

Table 1. Features of *lensmodel*

Feature	Capability
Point images	Position, flux, and time delay data. Image plane or source plane χ^2. Can handle multiple sources.
Arcs	Curve fitting.
Rings	Ring fitting with elliptical sources.
General data	Allows a floating registration between different types of data (e.g., radio and optical).
Mass models	Includes numerous circular and elliptical models, and allows arbitrary combinations of them.
Parameters	Gives full control over which parameters vary. Can produce arbitrary parameter surveys. Allows external constraints on parameter ranges. Allows relations between parameters to be imposed. Can use linear parameters and constraints.
Cosmology	Allows arbitrary values of Ω_M, Ω_Λ, and H_0. Can determine H_0 from time delays.
Output	Goodness of fit. Properties of mass model. Critical curves and caustics. Plots of time delay surface, potential, etc.

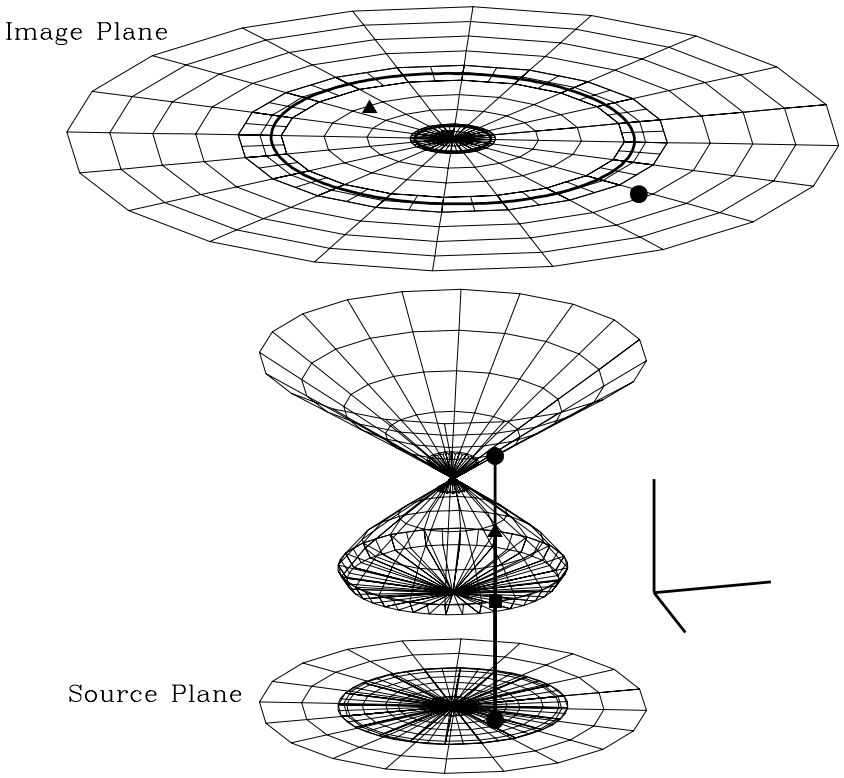

Fig. 1.— The geometry of the lens mapping for a nearly circular lens model. The top figure shows the tiling of the image plane. The middle figure shows how the image plane is distorted by the lens mapping. Height has been added to make the folding apparent. Projecting down through the distorted surface gives the tiling of the source plane (bottom), which overlaps itself to form multiply-imaged regions. The heavy curves in the image plane show the critical curves. A sample source is shown, together with the points where it intersects the distorted surface and where those points end up in the image plane. The central image is hard to see because the tiling is dense.

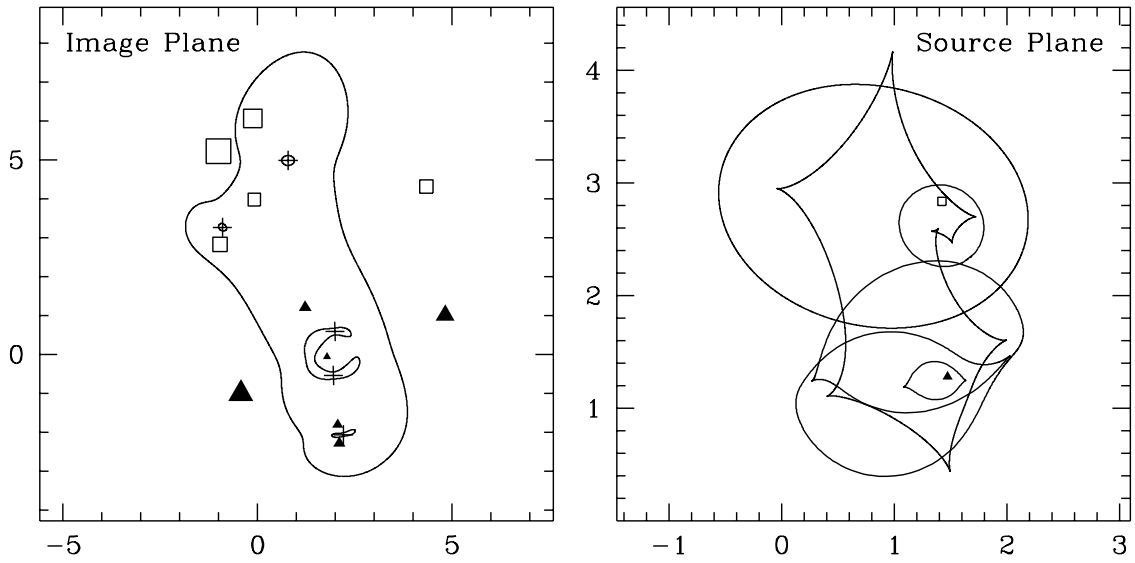

Fig. 2.— Sample critical curves and image configurations for a set of five random galaxies, computed with the tiling algorithm. In the image plane, the curves show the critical curves, the crosses mark the five galaxies, and the squares and triangles denote sample images. In the source plane, the curves show the caustics, and the square and triangle indicate the two sample sources. The square has seven images, of which five are bright and two are trapped in galaxy cores. The triangle has nine images, of which six are bright and three are trapped in galaxy cores. In the image plane, the size of each point is proportional to the magnification of the image. The axis scale is nominally arcseconds but is in fact arbitrary, because this example is intended only to demonstrate the capabilities of the tiling algorithm.

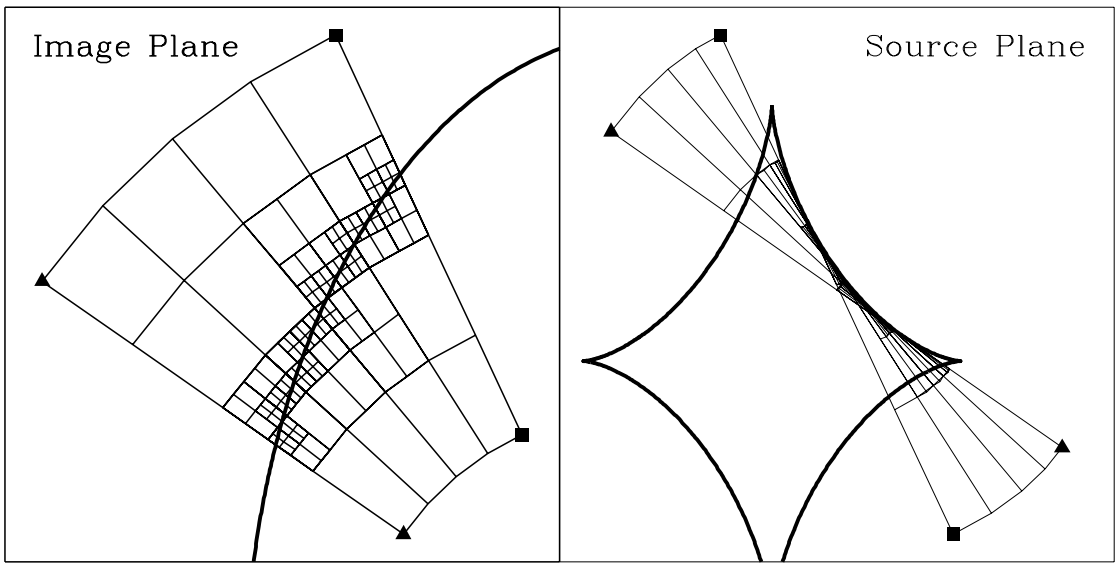

Fig. 3.— A sample tiling near a fold caustic, with three levels of sub-tiling. The left panel shows part of the image plane, and the right panel shows the corresponding part of the source plane. The heavy curves show the critical curve and caustic. The filled triangles and squares help illustrate the fold by matching vertices in the image and source planes to each other.

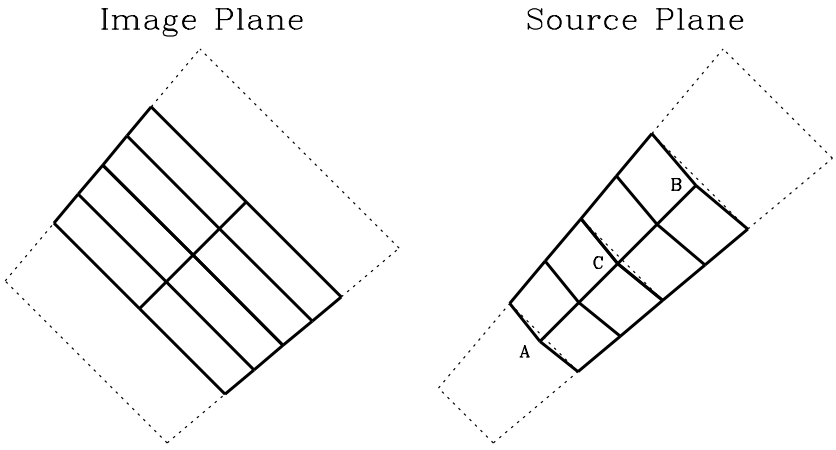

Fig. 4.— A sample tiling with four tiles, two of which have 2×2 sub-tiling; the light dotted lines show the main tiling, and the heavy solid lines show the sub-tiling. In the image plane the vertices of a daughter tile lie on the edges of its parent tile, but this may not be true in the source plane. The result can be an overlap (point A) or a gap (point B) in the source plane tiling. A gap or overlap appears only at places where the tiling changes resolution; point C shows that if adjacent tiles both have sub-tiles, the gap on one side is exactly offset by the overlap on the other.